Overview

Course: CS 141 – Computer Hardware

Course Level: Upper Level Undergraduate

Course Description: This course introduces fundamentals in designing and building modern information devices and systems that interface with the real world. It focuses on digital devices and systems, and it complements ENG-SCI 152, which focuses on devices and systems that use analog electronics. Topics include: combinational and sequential logic; computer architecture; machine code; and altogether the infrastructure and computational framework composing a MIPS processor. Consideration is given in design to interactions between hardware and software systems. Students will design application specific hardware for an embedded system.[1]

Module Topic: Hardware Backdoors and the Doctrine of Double Effect

Module Author: Zachary Gabor

Semesters Taught: Spring 2021

Tags: Hardware [CS], Privacy [both], intentions [phil], doctrine of double effect [phil]

Module Overview: Students are presented with examples of hardware backdoors and asked to think about both the benefits they may provide and the vulnerabilities they may create. The lesson discusses the doctrine of double effect as a tool for scrutinizing the acceptability of creating these risks in exchange for these benefits.

Connection to Course Material: Computer hardware design involves navigating different kinds of security vulnerabilities than those present in the design of software. There is, however, a basic ethical question which applies in both cases: when is it acceptable to expose users to a security vulnerability? Hardware backdoors, communication channels which operate below the level of a software OS, pose their own versions of this question.

The initial inspiration for the topic were the Specter and Meltdown hardware vulnerabilities. Though they were not the case studies used in this module, they might be useful for future iterations.

Goals

Module Goals:

1. Familiarize students with the Doctrine of Double effect along with auxiliary tools for determining how it applies to cases.
2. Evaluate both the utility and limit of the DDE in making moral assessments.
3. Apply this reasoning to real cases in which technical benefits and security concerns conflict.

Key Philosophical Questions:

1. When is it acceptable to do something which you know poses risks to others in exchange for benefits?
2. Is there a moral difference between intended and merely foreseen effects of one’s actions? Why or why not?
3. What are some possible issues and exceptions to the Doctrine of Double Effect as a guide to permissibility?

This module introduces students to the Doctrine of Double Effect and demonstrates how to utilize this framework effectively in applied ethical reasoning

[1] my.harvard

Materials

Key Philosophical Concepts:	- Doctrine of Double Effect - Intention - Means vs Side Effects - The “Why/How” Test in Intentions - The “Gerrymandering” objection to the DDE	As noted above, the goal of this module is to present a specific tool and demonstrate how to use it. Crucial to this latter goal is providing a tractable rule of thumb for distinguishing between consequences intended as means and those merely foreseen as side effects. For this purpose, the module covers the “why/how test” for making this distinction. The idea is that by asking “how do you mean to do that?” you can elicit an agent’s means, and by asking “why do you mean to do that?” you can elicit an agent’s ends.
Assigned Readings:	- Nienh-hê Hsieh and Rosemarie Monge “Recovering the Logic of Double Effect for Business: Intention, Permissibility, and Impermissible Harms”, pp. 1-10	This reading is an accessible introduction to the Doctrine of Double Effect and the Why/How test. The remainder of the paper may also be assigned in order to familiarize students with some common objections to the Doctrine of Double Effect as well as some examples of its application.

Implementation

Class Agenda:	1. Introduce the ethical challenges surrounding backdoors with two case studies: INTEL AMT and the NSA-designed Clipper Chip. 2. Introduce the Doctrine of Double Effect 3. Explain the How/Why Test as a tool for eliciting an agent’s means and ends 4. Discussion: objections and limitations of the DDE 5. Activity: applying the DDE to Intel AMT and the Clipper Chip	Intel AMT is a hardware product that allows corporate IT departments to perform various operations remotely on machines, some of which do not require the machine to be running an OS or even to be powered on. The security vulnerabilities it poses have been exploited by cybercriminals. The Clipper Chip was a digital wiretapping device devised by the NSA in the 1990s. Several security vulnerabilities were found to afflict it shortly after its development, and it never achieved widespread use. One way in which the Embedded EthiCS TA may motivate the DDE is to discuss the ways in which a straightforward cost-benefit

		analysis may be limited as a tool for distinguishing when it is permissible to cause harms in exchange for benefits and then to introduce the Doctrine of Double Effect and the How/Why test as tools to distinguish between permissible and impermissible action.
Sample Class Activity:	Students are asked to discuss in groups whether each case meets each of the criteria for permissibility according to the DDE: they are first asked to enumerate the harms and benefits which would accrue if the technologies are built and implemented, and to divide the harms into those which were means to the benefits and those which were mere side effects. They are then asked to determine whether the harms are proportional to the benefits and whether the creation of the technology would be an intrinsically good or neutral project. On the basis of these assessments, students are asked whether building the technology is permissible according to the DDE, and (perhaps independently of the verdict of the DDE) whether these technologies should be built.	Students are encouraged to understand the DDE as a heuristic to sort cases into those deserving more or less scrutiny. Consequently, it makes sense to ask students, in light of the DDE analysis they have undertaken, whether they think the technology should or should not be built.
Module Assignment:	The follow-up assignment consists of a final exam question. It includes a diagram of the "connected loop" version of the trolley problem, which the DDE arguably pronounces, counterintuitively, to be impermissible, and the following prompt: "In class we saw that, according to the Doctrine of Double Effect, it is allowable to switch the trolley in the standard setup of the trolley problem, killing the one to save the five. If we modify the setup so that the branches where the five and the one are stuck form a loop, as in the diagram below, is it still allowable to switch to the one to save the five according to the Doctrine of Double Effect? (Assume that running over one or more people will derail the trolley, so switching to the one will definitely save the five). Why or why not?"	Because the module is focused around a specific tool, a followup assignment prompting students to focus on an intricate, plausibly problematic application of that tool serves the module's skill-building purposes.
Lessons Learned:	One surprising result of this module was that in discussion, students concluded that Intel AMT, a basically uncontroversial technology in widespread use, should not be built. In future iterations, it might be a good idea either to prime students against expecting that conducting a DDE analysis will necessarily support the most	

	restrictive option available (as such an expectation might explain why students reached the radical conclusion they did) or to devote some time in subsequent discussion to underlining the surprisingness of this conclusion.	